



HAL
open science

Impact of point cloud normals compression on objective quality assessment

Amar Tious, Toinon Vigier, Vincent Ricordel

► **To cite this version:**

Amar Tious, Toinon Vigier, Vincent Ricordel. Impact of point cloud normals compression on objective quality assessment. 2024 32nd European Signal Processing Conference (EUSIPCO), Aug 2024, Lyon, France. pp.531-535, 10.23919/EUSIPCO63174.2024.10715168 . hal-04807718

HAL Id: hal-04807718

<https://hal.science/hal-04807718v1>

Submitted on 27 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Impact of point cloud normals compression on objective quality assessment

Amar Tious, Toinon Vigier, Vincent Ricordel

Nantes Université, École Centrale Nantes, CNRS, LS2N, UMR 6004, F-44000 Nantes, France

amar.tious@univ-nantes.fr

Abstract—With the emergence of immersive technologies and new use cases, point cloud (PC) compression and quality assessment have become very active research fields, which led to tremendous progress and standardization challenges in the last 5 years. But these methods only propose to compress the PC’s geometry and color properties, and very few studies also consider other attributes that could improve the rendering of point clouds in 3D scenes. In this paper, we focus on the surface normal properties that allow surface shading of the object, and add these attributes too in the video-based point cloud compression (V-PCC) pipeline considering two approaches: either by using attribute coding to transmit the normals, or by recomputing them from decoded PC geometry. We compare these approaches with objective metrics that take distortion of the normals and shaded rendering into account, and show the strategic advantages of normal attribute compression. Finally we discuss the challenges that point cloud compression brings up in order to optimize the physically-based rendering.

Index Terms—Point cloud compression, Objective quality assessment, Physically-based rendering, Point Cloud Quality

I. INTRODUCTION

Volumetric data permit the representation of real-life scenes and objects in 3D space. They can be captured by various acquisition methods that fit to different use cases, such as LiDAR scans for large scenes or multi-camera studios for objects and people. Point cloud describes volumetric content as a set of points characterized by 3D coordinates and various attributes that describe the object texture and surface. It’s an efficient alternative to 3D meshes because it’s easier to store and computationally cheaper. Thanks to the late progress in the field of eXtended Reality (XR), point clouds have become relevant to plenty of new use cases and challenges including immersive 3D video streaming, architecture & interior decoration, virtual museums and telepresence. These application fields require efficient data compression for storage and transmission purposes, rendering methods and quality assessment (QA). For visual QA and point cloud compression (PCC), the experts group projects MPEG-I [1] and JPEG Pleno [2] have defined standards and test conditions. Precisely, subjective QA tests aim at scoring the visual quality of distorted point clouds directly perceived by observers. In parallel, objective QA metrics [3], [4] are developed in order to computationally evaluate decoded point cloud quality and aim at producing scores highly correlated with human judgment. In each case, the context of the final use case should be considered [5].

Yet so far, most QA studies for PCC consider only point coordinates and color attributes, and never consider

advanced rendering methods with shading and dynamic surfaces. Rossoni et al. [6] proposed a PCC pipeline that takes into account attributes describing the surface and material properties of the object, in order to allow physically-based rendering (PBR), but no attempts of implementation for such a pipeline existed until now. Neither were there any QA studies on the impact that it would have on the distorted point cloud’s visual quality.

In this paper, we take a first step toward the visual improvement of decoded point cloud rendering by adding the coding of the point normals in video-based point cloud compression (V-PCC). Namely, in the V-PCC scheme, we code the PC geometry (3D point coordinates), and two attributes: the RGB colors and the normals, the coordinates of vectors normal to the points surface planes. To prove the advantage of our V-PCC scheme, we compare to the case where normals are not transmitted and are instead computed from the distorted geometry on the decoder side. At this level, we identify existing objective metrics that account for normals to assess the quality of distorted content. The research questions we investigate in this paper are then the following:

- In the PCC pipeline, is it better to code the normals than to estimate them from the decoded geometry?
- Considering the V-PCC pipeline and PBR, how do current objective QA metrics assess the quality of rendered PC?

II. EXPERIMENTAL METHODS

A. V-PCC and normals encoding

We used the MPEG-I V-PCC standard [1], which has been designed for volumetric video compression, and its reference software *tmc2* [7]. The standard first converts by projection the PC data into 2D patch maps corresponding to the point attributes and coordinates. The resulting texture, occupation and depth maps are then encoded by a selected video codec (HEVC in *tmc2*). In our case, we consider that the original PC also has surface normals as attributes. In order to add the normals in the transmission pipeline, we considered two approaches: either we encode the normals like the colors, or we prune them prior to encoding and re-estimate them from the decoded points geometry.

1) *Encoding of the normals*: The V-PCC encoding process fits to any attribute that can be represented as a 2D projection map. The normals (n_x, n_y, n_z) being a 3-dimensional attribute, can be projected onto a normal map and encoded using the




Fig. 1. General framework prototype with normals attribute coding. Normals are set and transmitted as a normal texture, allowing to maintain the orientation and direction of the reference surface.




Fig. 2. General framework with computation of normals. The original normals are pruned prior to the encoding, and re-estimated after transmission to allow a shaded rendering.

same process as for the RGB color attributes. In our case, we need first to quantize the normal coordinates from floating values between $[-1 : 1]$ to integer values between $[0 : 255]$, which induces a negligible loss. Fig.1 describes the framework used to encode the point cloud normals: they are quantized and then set and processed as the points colors in order to produce a "3D normal texture". Then, the normal texture and the PC with its original color attributes are encoded using V-PCC. After decoding, the processed normal maps are converted back to normalized vector coordinates and assigned as normal attributes to the decoded colored PC again.

This process is only an early strategy that fits our research. In the context of compression for transmission, consideration for normals will have to be implemented in the V-PCC pipeline to avoid encoding the geometry twice.

2) *Recomputing normals after decoding*: Currently, for PCC use cases where the normals are needed for surface description or shading after decoding, they are recomputed using the following method: the surface plane for each point is estimated from its k-nearest neighbor (KNN) points, and the coordinates of the normal vector of this plane are set as point attributes. This approach, depending only on the point coordinates, lacks information about the original orientation of the normals and the smoothness of the surface. It performs poorly for too sparse or too dense point clouds and requires human control to guarantee satisfying results. To improve the normals estimation with no human control, we add few steps which correct simple artifacts: after a first estimation of the

normals, we use Poisson Surface Reconstruction [8] and re-orient the normals according to the constructed surface. We then apply a smoothing filter to the normals and assign them to the decoded geometry. Fig.2 illustrates the corresponding framework.

B. Objective QA metrics for normal distortion

There are two main categories of objective QA metrics for point clouds: point-based metrics that determine a quality score from the raw point data, and image-based metrics that first project the 3D data into multiple 2D views and then evaluate 2D image quality.

1) *Point-based metrics*: They use the point cloud's geometry and attributes as inputs to compute a quality score. Until now in almost all subjective QA tests, point clouds are unlit and rendered by using only the point coordinates and colors. Similarly, the best performing point-based metrics consider these same conditions in their score computation. So in our case which takes processed normals into account and uses them in the PC rendering, most of the state-of-the-art tools don't fit. There are however some metrics that use the normals as geometrical clues to evaluate the surface distortion of decoded point clouds, or assess the quality of any given attribute. The metrics from JPEG Pleno's [2] common test conditions (ctc) for PCC are standards which are the most frequently used in objective QA benchmarks and applications. They can easily be adapted to our needs: among them, the plane-to-plane measure (AngSIM [9]) assesses the point cloud's geometry quality with the normal orientation. Originally meant to be computed on estimated normals for both the reference and evaluated stimuli, we can use it to rate the quality of distorted normals. The same goes for the point-to-plane measure (PSNRD2 or p2plane [10]) which is the dot product between the geometric error of point coordinates and the normals. At last, the point-to-point (PSNRD1 [10]) attribute measure can be adapted to normal vector coordinates rather than YUV colors.

2) *Image-based metrics*: With them, it is possible to not only consider the 2D-projected raw data, but also distortions and characteristics induced by the rendering parameters. Considering videos or images of point clouds rendered in a 3D environment, state-of-the-art image quality metrics (IQMs) can be used to evaluate the stimuli as it would be displayed and perceived by observers in a subjective test. Previous studies [11], [12] show that some IQMs are among the best performing metrics for point cloud QA.

In order to consider normals with this method, we used videos of shaded point clouds obtained from the point of view of a virtual camera rotating around the barycenter of the 3D object. As representation of the surface of the PC, the normals can be used in the rendering process in order to compute the lighting of points and therefore allow the shading of the objects. QA studies for 3D graphics [13] have shown the impact of lighting in geometry perception. Because it is a popular platform for development of XR applications and as we suggested in [14], we used the Unity Game Engine to

generate these videos. We rendered the stimuli, we designed a PBR shader allowing us to adjust the points size and make the stimuli surfaces look watertight. The Unity 3D scene we synthesized had no skybox and environment, just a plain medium gray background, a directional light facing the object from 30° above as advised by [15] to simulate a sunlight, and a 25% intensity ambient light to make the shadows noticeable but not too masking. The camera rotation around the object was the same as used in [11]: 420° degrees around the y -axis (-30° to 390°) and 60° around the x -axis (0° to 60°). The length of the videos was 10 seconds as recommended by JPEG Pleno. Fig. 4 shows point clouds rendered as described, and reveals first perceivable artifacts and differences that would impact observer ratings in subjective QA.

III. EXPERIMENTAL RESULTS



Fig. 3. Screenshots of the selected PC references, rendered with no shading.

A. Data Preparation

We selected references from BASICS [11], being the most voluminous and diverse PC dataset (75 sources, 1500 processed PC, 4 semantic categories) available with subjective test results (observer ratings and MOS). We selected a subset containing only content from 3 semantic categories (*humans*, *inanimate objects*, *buildings*) that have ground-truth normals and that represent 3D data acquired by relevant methods (LiDAR for buildings, multi-camera capture for others, but no computer-generated graphics). We then chose 2 models per category based on their original visual quality and density. The 6 selected PC references are shown in Fig. 3.

For each source point cloud, we generated new distorted versions with V-PCC considering different geometric and attribute compression parameters. Since the color distortion effect isn't in our interest in this study, we decided to use a lossless compression for the color attributes of all distorted PC versions. We selected 3 different bitrates from the JPEG Pleno Point Cloud coding *etc* [2] for the geometry compression ($g1 = 0.03bpp$, $g3 = 0.05bpp$, $g5 = 0.12bpp$). For each of these geometries, we generated 4 normal distortion levels ($n1 = 0.07bpp$, $n3 = 0.15bpp$, $n5 = 0.36bpp$, nE). The nE level corresponds to normals computed from the distorted geometries. Fig. 4 displays for one PC reference for each semantic category its corresponding two most distorted versions,

the most compressed geometry with the most compressed normals ($g1-n1$), and with estimated normals ($g1-nE$).



Fig. 4. Examples from each semantic category of a reference point cloud and two highly distorted versions ($g1-n1$ and $g1-nE$) rendered in Unity with a directional light. The surface artifacts of estimated normals (nE) are revealed by the virtual lighting.

B. Objective QA

From a selection of full-reference point-based metrics (see II-B1) and IQMs, we have objectively assessed the quality of the distorted point clouds in order to compare estimated versus encoded normals cases. By scoring the quality of the rotating camera videos of shaded point clouds, we also objectively assessed the normal distortion effect in the context of rendering, because the normals represent the object surface from the PBR shader's perspective. Exactly, we used 4 well known IQMs: the PSNR, the SSIM [16], the perceptual metric VIF [17], and the multi-method VMAF [18]. Quality scores are computed for each frame of the videos and combined into an overall score by an arithmetic mean [18].

In order to observe the impact of shaded rendering on objective QA, we also generated additional rotating camera videos for each different geometry but without directional sunlight in the 3D scene. Instead, a 100% intensity ambient light is used, such as the PC is rendered without consideration of the normals. In the tables, these videos are indexed as $n0$. In order to score the quality of the $n0$ videos, the image-based QA approach is used, because the distortion comes only from the rendering.

Table I presents the average metric scores for each distortion parameter. Table II shows the average metric scores considering the $n3$ and nE distortions levels of each semantic category of content. We computed the average scores displayed in the table using all geometry compression levels for the given normal distortions. The "median" level $n3$ is displayed to represent the case of normal compression.

IV. DISCUSSION

In Table I, considering the metrics individually, the results show that some metrics are less affected by some distortions. The $p2plane$ does almost not vary with the bitrate of

TABLE I

AVERAGE OBJECTIVE QUALITY SCORES FOR EACH DISTORTION TYPE: INCREASING COMPRESSION LEVELS FOR NORMALS ($n1, n3, n5$) AND FOR GEOMETRY ($g1, g3, g5$), ESTIMATED NORMALS (nE), AND NO LIGHTING ($n0$). THE *max* COLUMN DISPLAYS THE POSSIBLE MAXIMUM OF EACH METRIC.

Category	Metric	$g1$					$g3$					$g5$					<i>max</i>
		$n1$	$n3$	$n5$	nE	$n0$	$n1$	$n3$	$n5$	nE	$n0$	$n1$	$n3$	$n5$	nE	$n0$	
Point-Based	Attr-PSNR	25.7	27.2	28.2	14.8	-	26.2	28.0	29.2	16.2	-	26.2	28.1	29.5	20.3	-	∞
	p2plane	68.6	68.6	68.6	65.8	-	72.1	72.1	72.1	69.6	-	75.4	75.4	75.4	72.4	-	∞
	AngSIM	.917	.933	.943	.857	-	.918	.935	.945	.886	-	.922	.939	.949	.909	-	1
IQM	PSNR	37.7	37.8	37.8	36.5	36.9	38.9	39.1	39.1	37.6	38.1	39.7	39.9	40.0	39.4	38.7	∞
	SSIM	.967	.967	.968	.966	.958	.972	.972	.973	.971	.964	.973	.974	.974	.973	.965	1
	VIF ₃	.872	.879	.881	.839	.901	.901	.908	.911	.880	.926	.916	.925	.928	.907	.941	1
	VMAF	52.0	53.4	53.9	45.6	52.5	61.1	62.6	63.2	57.1	61.3	66.2	68.0	68.7	64.4	65.2	100

TABLE II

AVERAGE OBJECTIVE QUALITY SCORES FOR THE $n3$ AND nE DISTORTION LEVELS AND CONSIDERING THE SEMANTIC CONTENT CATEGORY.

Metric	All		Human		Object		Building	
	$n3$	nE	$n3$	nE	$n3$	nE	$n3$	nE
Attr-PSNR	27.7	17.1	34.9	25.3	28.8	18.9	19.5	7.1
p2plane	72.0	69.3	73.3	73.2	72.7	72.6	70.0	61.9
AngSIM	.936	.884	.972	.936	.968	.924	.868	.792
PSNR	38.9	37.9	43.9	41.9	38.6	37.4	34.3	34.2
SSIM	.971	.970	.993	.992	.970	.969	.950	.950
VIF ₃	.904	.876	.969	.942	.927	.879	.816	.806
VMAF	61.4	55.7	78.1	73.0	66.3	56.2	39.6	38.7

compressed normals, and the image-based SSIM only shows a slight difference in quality between decoded normals and estimated normals. Otherwise, the scores generally follow the quality ranking: $g5 > g3 > g1$ and $n5 > n3 > n1 > n0 > nE$. For all metrics, the quality for nE (estimated normals) is also judged noticeably better for the $g5$ geometry than for $g3$ and $g1$. The image-based metrics are also more impacted by distortion of the geometry than of the normals. This is a satisfying observation, considering that the normals, being 3-dimensional attributes, are more expensive to transmit than geometry. High quality point coordinates with low quality normals could be better than low quality coordinates with high quality normals, both in terms of bandwidth and visual quality.

As all the metric scores show it, even highly distorted normals ($n1$) have a better quality than estimated normals (nE). From point-based results, it means that surfaces rendered with encoded normals look more faithful to the reference surfaces. The image-based results suggest that, in the context of shaded rendering, compressing the normals as attributes can mask the geometry coding artifacts. While on the contrary, estimated normals which are computed from the decoded point coordinates, are more revealing of the rough "patchy" distorted geometry induced by V-PCC. In comparison to unlit rendering, it reveals that shading with estimated normals has a negative impact on the visual quality, as observed from subjective test results too in [19]. However in the latter paper, authors only assess the quality with estimated normals and not coded ones. According to our results, decoded normals can bring an increase in visual quality when used for surface shading. This encourages the development of point cloud rendering methods.

In Table II, a noticeable difference on point clouds representing *buildings* appears: their quality scores are significantly lower than for other semantic categories. Indeed, geometry coding, attribute coding and normal estimation perform poorly for this type of content. These observations coincide with previous objective studies [11], [20] and can be interpreted in two ways: either other compression methods than V-PCC are more efficient for this type of content, or the modern QA metrics for point clouds perform poorly for data acquired with LiDARs.

Encoding the normals has a cost in transmission bitrate, whereas estimating them has not. But as our experimental measures suggest, normals and coordinates even with the heaviest distortion ($g1-n1$) still offer an equivalent if not better visual accuracy than estimating the normals from very low geometry distortion ($g5-nE$) or than unlit rendering ($n0$). So transmitting normals, even at a low bitrate, possibly permits to save bits for a more economic transmission for the geometry. Moreover, using them in an application with shaded rendering or PBR can improve the visual quality in comparison with unlit surfaces. To some extent, since transmitted normals describe a surface closer to the original, their orientation could be used to fix strongly distorted geometry and correct still noticeable artifacts. Besides, each method, compression or estimation, has its own computational cost that could be investigated and compared.

The noticeable difference in visual quality between $n1$ and nE on Fig. 4 are the same that we identified through objective QA, and could be accused in subjective tests as well. Even though our experimental metric scores show the relevance of normal compression, the visual aspect of this strategy should be validated through subjective QA. This would provide us with ground-truth quality scores to measure and compare the performance of the used metrics and their correlation with the human visual system and subjects' opinions. Subjective QA could also help in identifying needs and thresholds to develop new objective QA tools, which is another need brought up by this paper. Considering the weight of normals and shadows in the computation of a visual score could enhance the performance of objective QA. Similarly the visual attention data could help in weighing regions of interest for the rating [21], [22].

This paper only focuses on the case of normals distortion

and its interaction with geometry distortion, so it didn't consider color compression. In other objective QA studies and in realistic and practical scenarios, the points' colors are compressed as well [2], [11] and their distortions have shown heavy weight in the resulting visual quality. Further studies considering both lossy normals and color attribute coding have to be considered for a more realistic experiment setting. Other attributes such as the points transparency, reflectance, emission, or surface metalness and roughness (as suggested in [6]) could be encoded or estimated, and would have to be considered for future QA studies.

V. CONCLUSION

In this paper, we made a step forward a more complete compression scheme for point cloud attributes in order to enhance the rendering of decoded PCs. We conducted an objective QA study comparing different approaches for considering the normals in a compression scenario. By encoding these normals as attributes in the V-PCC pipeline, we showed that there is potential gain in terms of quality and even bandwidth rather than recomputing them from the decoded geometry. We showed the advantages of surface shading rather than simple unlit rendering. Recomputing the normals on distorted geometry reveals artifacts and lowers the visual quality of the rendered point cloud. Encoded normals, even at low bitrate, enhance the visual aspect of rendered point clouds, and should help in saving bits for geometry compression. Further studies are needed to definitively validate our first results and conjunctions. Subjective QA tests should be run to collect ground-truth quality reference ratings and assess the performance of point cloud QA metrics in this scenario. Besides, upgrading the PCC codecs to efficiently consider new attributes, and adapting objective metrics to fit to new compression scenarios are promising research topics.

REFERENCES

- [1] S. Schwarz *et al.*, "Emerging mpeg standards for point cloud compression," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 9, no. 1, pp. 133–148, 2018.
- [2] S. Perry, "Jpeg pleno point cloud common test conditions v3.6," *ISO/IEC JTC1/SC29/WG1 N*, vol. 86044, Apr. 2021.
- [3] G. Meynet, Y. Nehmé, J. Digne, and G. Lavoué, "PCQM: A Full-Reference Quality Metric for Colored 3D Point Clouds," in *2020 Twelfth International Conference on Quality of Multimedia Experience (QoMEX)*, pp. 1–6, May 2020. ISSN: 2472-7814.
- [4] E. Alexiou and T. Ebrahimi, "Towards a Point Cloud Structural Similarity Metric," in *2020 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pp. 1–6, July 2020.
- [5] E. Alexiou and T. Ebrahimi, "Impact of visualisation strategy for subjective quality assessment of point clouds," in *2018 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, pp. 1–6, IEEE, 2018.
- [6] M. Rossoni, M. Pozzi, G. Colombo, M. Gribaudo, and P. Piazzolla, "Physically-based Rendering of Animated Point Clouds for eXtended Reality," *Journal of Computing and Information Science in Engineering*, pp. 1–9, Sept. 2023.
- [7] "GitHub - MPEGGroup/mpeg-pcc-tmc2: Video codec based point cloud compression (V-PCC) test model."
- [8] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson surface reconstruction," in *Proceedings of the fourth Eurographics symposium on Geometry processing*, SGP '06, (Goslar, DEU), pp. 61–70, Eurographics Association, 2006.
- [9] E. Alexiou and T. Ebrahimi, "Point Cloud Quality Assessment Metric Based on Angular Similarity," in *2018 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6, July 2018. ISSN: 1945-788X.
- [10] D. Tian, H. Ochimizu, C. Feng, R. Cohen, and A. Vetro, "Geometric distortion metrics for point cloud compression," in *2017 IEEE International Conference on Image Processing (ICIP)*, pp. 3460–3464, Sept. 2017. ISSN: 2381-8549.
- [11] A. Ak *et al.*, "Basics: Broad quality assessment of static point clouds in a compression scenario," *IEEE Transactions on Multimedia*, 2024.
- [12] Q. Liu, H. Su, Z. Duanmu, W. Liu, and Z. Wang, "Perceptual Quality Assessment of Colored 3D Point Clouds," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, pp. 3642–3655, Aug. 2022. Conference Name: IEEE Transactions on Visualization and Computer Graphics.
- [13] K. Vanhoey, B. Sauvage, P. Kraemer, and G. Lavoué, "Visual Quality Assessment of 3D Models: On the Influence of Light-Material Interaction," *ACM Transactions on Applied Perception*, vol. 15, pp. 5:1–5:18, Oct. 2017.
- [14] A. Tious, T. Vigier, and V. Ricordel, "Physically-based Lighting of 3D Point Clouds for Quality Assessment," in *Proceedings of the 2023 ACM International Conference on Interactive Media Experiences, IMX '23*, (New York, NY, USA), pp. 423–426, Association for Computing Machinery, 2023.
- [15] J. P. O'Shea, M. S. Banks, and M. Agrawala, "The assumed light direction for perceiving shape from shading," in *Proceedings of the 5th symposium on Applied perception in graphics and visualization, APGV '08*, (New York, NY, USA), pp. 135–142, Association for Computing Machinery, 2008.
- [16] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [17] H. R. Sheikh and A. C. Bovik, "Image information and visual quality," *IEEE Transactions on image processing*, vol. 15, pp. 430–444, Jan. 2006.
- [18] Z. Li, A. Aaron, I. Katsavounidis, A. Moorthy, M. Manohara, *et al.*, "Toward a practical perceptual video quality metric," *The Netflix Tech Blog*, vol. 6, no. 2, p. 2, 2016.
- [19] A. Javaheri, C. Brites, F. Pereira, and J. Ascenso, "Point cloud rendering after coding: Impacts on subjective and objective quality," *IEEE Transactions on Multimedia*, vol. 23, pp. 4049–4064, 2021.
- [20] D. Lazzarotto, E. Alexiou, and T. Ebrahimi, "Benchmarking of objective quality metrics for point cloud compression," in *2021 IEEE 23rd International Workshop on Multimedia Signal Processing (MMSP)*, pp. 1–6, Oct. 2021. ISSN: 2473-3628.
- [21] Z. Wang, Y. Zhang, Q. Yang, Y. Xu, J. Sun, and S. Liu, "Point cloud quality assessment using 3d saliency maps," in *2023 IEEE International Conference on Visual Communications and Image Processing (VCIP)*, pp. 1–5, IEEE, 2023.
- [22] J. Gutiérrez, G. Dandyeyeva, M. D. Magro, C. Cortés, M. Brizzi, M. Carli, and F. Battisti, "Subjective evaluation of dynamic point clouds: Impact of compression and exploration behavior," in *2023 31st European Signal Processing Conference (EUSIPCO)*, pp. 675–679, 2023.